








***Misconceptions and Old Wives' Tales -
On DB2 Database Maintenance and Recovery***




SOFTWARE ENGINEERING

Overview

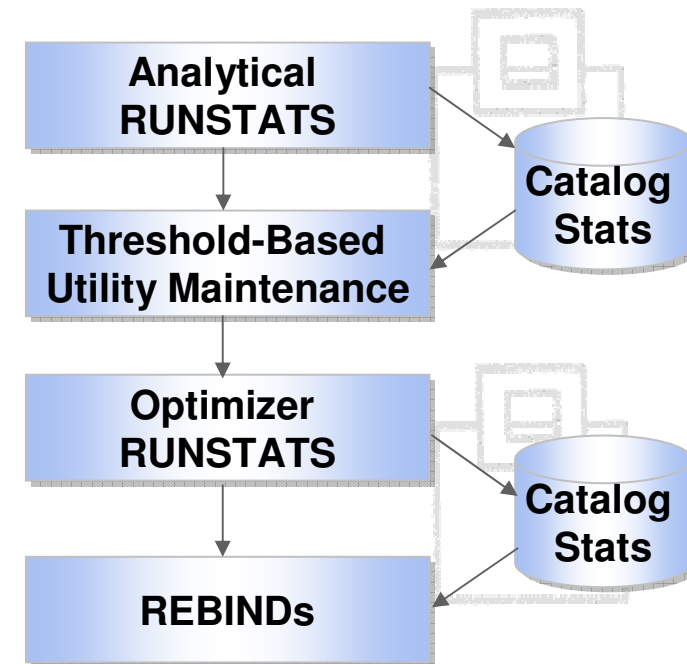
- Analytical RUNSTATS vs. Optimizer RUNSTATS
- Real-time statistics
- Backup considerations
 - Migrating to tape
 - Using VTS
 - Active and archive logs
 - Flashcopy
- Recovery time objectives (RTOs)
- Relevant ZPARMS for recovery
- Coupling facility performance vs. its recovery
- VSAM clusters for multi-volume data sets




Lies, Damned Lies, and Statistics




- DB2 statistics are the lifeblood of database health
- Historically two types of RUNSTATS:
 - Analytical RUNSTATS to analyze maintenance needs
 - Optimizer RUNSTATS to improve performance



SOFTWARE ENGINEERING

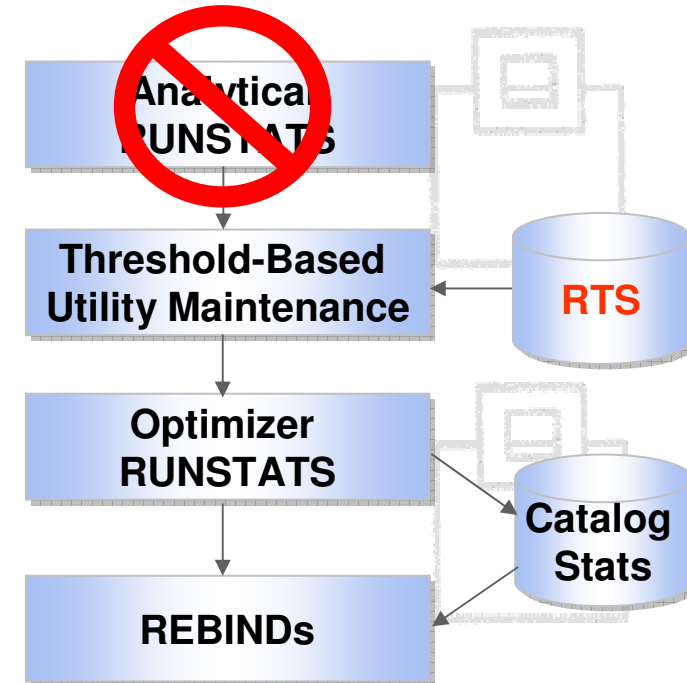


Lies, Damned Lies, and Statistics




Misconception: Using RTS costs CPU

- Since V7 there are two types of statistics:
 - Catalog statistics
 - Real-time statistics
- Stop using catalog statistics for analytical RUNSTATS
- RTS are collected even if you don't use them
- Use real-time statistics for accurate statistics




→ Save CPU! Eliminate analytical RUNSTATS

SOFTWARE ENGINEERING




Lies, Damned Lies, and Statistics



RTS pertaining to backups:


- **COPYLASTTIME**

The timestamp of the last full or incremental image copy on the table space or partition.




- **COPYUPDATEDPAGES**

The number of distinct pages that have been updated since the last COPY.



- **COPYCHANGES**

The number of insert, delete, and update operations since the last COPY.



- **COPYUPDATELRSN**


The LRSN or RBA of the first update after the last COPY.

- **COPYUPDATETIME**

The timestamp of the first update after the last COPY.



SOFTWARE ENGINEERING

Lies, Damned Lies, and Statistics




What's new with RTS in DB2 9

- The RTS tables are moved to the DB2 catalog
- Some new statistics are added
e.g., index usage information



**New in
DB2 9**



SOFTWARE ENGINEERING


All Backups are There!

- You've got your statistics straight
- And, all backups are there!
- But, do the backups fully serve their purpose?






SOFTWARE ENGINEERING

Migrating To Tape




Misconception: Free up disk space ASAP

- Don't migrate to tape too early
- Which is more important, freeing up a little disk space or time to recover ?
- Consider keeping the most current backups on DASD
- To save space, limit dual backups to critical copies:
i.e., after REORG or LOAD LOG NO






SOFTWARE ENGINEERING

Virtual Tape Storage



Misconception: VTS is as good as DASD





- Do not overestimate the speed of virtual tape story
- Always keep mount and recall times in mind
e.g., for small objects these are often bad in relation to recovery time



SOFTWARE ENGINEERING


Oops! The archive logs are on tape

- Then, a parallel recovery is not possible
 - When backing up archive logs to tape or VTS:
 - Do NOT use 28672 block size (normally optimal for tape)
 - Use 24576, which is preferred for DASD
- ZPARM BLKSIZE *24576 default since V8*
- General ROTs:
 - Keep active logs on DASD 24 hours
 - Keep archive logs on DASD 48 hours



SOFTWARE ENGINEERING




Always Use Dual Logging



Misconception: Single active logging is OK if mirroring is used


- Always use dual logging to assure data consistency

→ ZPARMS TWOACTV
TWOARCH



SOFTWARE ENGINEERING




Using FlashCopy



Misconception: „End of the image copy“


- Advantages:

- IBM' V8 BACKUP SYSTEM and RESTORE SYSTEM takes full advantage of FlashCopy
- Immensely fast
- Frees the CPU because the I/O subsystem does the work
- Using consistency groups, FlashCopy 2 supports incremental ICs at the volume level and over multi-volumes






SOFTWARE ENGINEERING

Using FlashCopy



Misconception: „End of the image copy“


- But not the end of image copy for everyone:
 - Expensive in terms of the hardware and system topology
 - Twice the number of disks required to flash a given pool
 - Offloading to tape fills up gigabytes of unchanged data
- FlashCopies are not registered anywhere except by the BACKUP SYSTEM DB2 command
- Flashes don't reset the the COPY PENDING status



→ Choose objects wisely for FlashCopy




SOFTWARE ENGINEERING

Copying Indexes




Misconception: Don't need COPY YES for indexes

- Many still fail to use COPY YES for indexes
- COPY YES allows image copy backups of indexes
- Allows the recovery using either REBUILD or RECOVER
- Reduces overall outage by allowing the recovery of both the data object and the index




SOFTWARE ENGINEERING

R U Really Recovery Ready?




Misconception: All backups are available and current


- Large shops are carefully aligning backup frequency to update rates using automation, thresholds, and monitoring
- What's wrong with this picture . . .




→ Does the recovery duration support the business needs?



SOFTWARE ENGINEERING





Recovery Duration



Misguided notions:



- Recovering small objects with minimal update rates will be fast
(Not if those few updates cover several archive logs!)
- Nightly backups of large objects with high updates rates are enough
(Nope. These require extensive, time-consuming, log applies to recover during the day.)




SOFTWARE ENGINEERING



Recovery Duration

DB2
ICs/IICs




TS-B
LOG NO
Event



	Pages	%Updt			RTO	Mount	Rec.
TS-A	5.000	2%	2	1	12:02	3 LOG	
TS-B	50.000	10%	0	2	n/a	n/a	n/a
TS-C	500.000	50%	0	1	3:37	1 IC	




SOFTWARE ENGINEERING

A, B, Z's of ZPARMS




Misconception: ZPARM defaults are dependable

- ZPARM defaults are not „one size fits all“
- Migrating to a new DB2 version does not apply new defaults






SOFTWARE ENGINEERING

Fast Log Apply




- ZPARM LOGAPSTG
- Best value is 100MB
- Default of 100MB changed in V8
- Never assign less than 100MB
- If DB2 needs less, it simply takes less






SOFTWARE ENGINEERING

Check Frequency




- ZPARM CHKFREQ
- Determines check point frequency in minutes or number of log records
- Default is 500,000 records
- Log intervals can miss checkpoints if logging rates vary a lot
- Consider changing this to minutes to maximize performance
Recommended 2 to 5 minutes






SOFTWARE ENGINEERING

Retention Lock Wait




Note: For data sharing only

- ZPARM RETLWAIT
- A multiplier on how long a member of a data-sharing group waits for locks on a resource held by another member that has failed
- Default is 0, which means the lock request is immediately rejected and the application gets a resource unavailable SQLCODE
- Should be set to 2






SOFTWARE ENGINEERING

Output Buffer




- ZPARM OUTBUFF
- Output buffer size used for writing active log data sets
- Default since DB2 V8 is 4000
- Popular health checks recommend 40000






SOFTWARE ENGINEERING

Coupling Facility (CF)




- A cause of blindness?

$\frac{20}{200}$	E	$\frac{200 \text{ FT}}{61 \text{ M}}$	1
$\frac{20}{100}$	F P	$\frac{100 \text{ FT}}{30.5 \text{ M}}$	2
$\frac{20}{70}$	T O Z	$\frac{70 \text{ FT}}{21.3 \text{ M}}$	3
$\frac{20}{50}$	L P E D	$\frac{50 \text{ FT}}{15.2 \text{ M}}$	4
$\frac{20}{40}$	P E C F D	$\frac{40 \text{ FT}}{12.2 \text{ M}}$	5
$\frac{20}{30}$	E D F C Z P	$\frac{30 \text{ FT}}{9.14 \text{ M}}$	6
$\frac{20}{25}$	F E L O P Z D	$\frac{25 \text{ FT}}{7.62 \text{ M}}$	7
$\frac{20}{20}$	D E F P O T E C	$\frac{20 \text{ FT}}{6.10 \text{ M}}$	8
$\frac{20}{15}$	L R F O D P C T	$\frac{16 \text{ FT}}{4.57 \text{ M}}$	9
$\frac{20}{13}$	F D P L T C H O	$\frac{13 \text{ FT}}{3.96 \text{ M}}$	10
$\frac{20}{10}$	F R O C C F P D	$\frac{10 \text{ FT}}{3.05 \text{ M}}$	11



SOFTWARE ENGINEERING




Coupling Facility (CF)



Misconception: Good CF performance indicates a fast recovery


If a member dies:

- All data sets are put into LPL or GRECP
- Lots of jobs with up to 99 –STA commands must be built and submitted
- Only a finely-tuned CF can withstand this without killing itself





SOFTWARE ENGINEERING

Coupling Facility (CF)




IBM's metric guidelines to enable speedy recovery of LPL or GRECP:

- "CF transfer time" less than 2,000 microseconds.
- "Number of messages rejected due to lack of message buffer" zero.
- "Sub channel busy" percentage less than 10%.
- "False lock contention" percentage less than 3%.
- "All path busy termination count" zero.
- GBP "cross invalidations due to directory reclaims" zero.
- GBP "asynchronous failed writes due to lack of storage" zero





Coupling Facility (CF)




Useful assembler macros:

- IXCQUERY requests information about resources the CF manages.
- IXLMG requests measurement data related to the use of a CF.
- IXLZSTR retrieves control information and structure data from a dump.




SOFTWARE ENGINEERING

Multi-Volume Data Sets




Misconception: Multi-volume data sets don't matter

- A DASD space shortage is a DASD space shortage is a DASD space shortage - multi-volume data sets alone are not going to help.
- Don't let DB2 table spaces spread out over an unforeseeable number of volumes - "facilitated" by SMS DATA CLASS parameters.
- Multi-volume data sets can soak up below-the-line storage and make I/O, especially OPEN and CLOSE, slow.
- Highly fragmented disks worsen problems with multi-volume data sets.




SOFTWARE ENGINEERING

Improving Your *Insight*




Tips:

- Use real-time statistics to eliminate analytical RUNSTATS.
- Don't migrate backups too early to tape.
- Limit dual copies to critical copies.
- Virtual tape storage is not as good as DASD.
- Use a block size of 24576 if you must place archive logs on tape.
- Always use dual logging.
- Choose objects for FlashCopy wisely.
- Use the COPY YES when creating or altering indexes.
- Consider recovery time objectives when creating backups.
- Know your ZPARMS well.
- Avoid multi-volume data sets.
- Check for a healthy coupling facility.



Next Steps & More Information

- **Ask the speaker** u.heinrich@seg.de
- **Are you ready for DB2 9 for z/OS ?**
Contact your local IBM representative or email
WW DB2 for z/OS Market Manager Surekha21@uk.ibm.com
- **Need More Information**
<http://www-01.ibm.com/software/data/db2/zos/family/>
<http://www-01.ibm.com/software/data/db2/library/>
- **DB2 for z/OS Whitepaper** 
[DB2 9 for z/OS Data On Demand](#)
- **IBM Redbooks** [http://www.redbooks.ibm.com/Latest Redbooks](http://www.redbooks.ibm.com/LatestRedbooks)
- **Join "The World of DB2 for z/OS !"** <http://db2forzos.ning.com/>
- **Twitter DB2:** <http://twitter.com/IBMDB2>

**DB2 9 for z/OS
is here**



SOFTWARE ENGINEERING